CLAIMS

WHAT IS CLAIMED IS:

1. A communications system, comprising:

a first queue pair (QP) associated with a first connection, the first QP comprising a first send queue (SQ) and being associated with a first limit value;

a second QP associated with a second connection, the second QP comprising a second SQ and being associated with a second limit value;

a general pool comprising a shared receive queue (SRQ), the SRQ being shared by the first QP and the second QP; and

a resource manager providing provisioning for the SRQ,

at least one of the resource manager, the first QP and the second QP managing the first limit value and the second limit value.

2. The communications system according to claim 1,

wherein the first QP manages the first limit value, and

wherein the second QP manages the second limit value.

3. The communications system according to claim 1, wherein the first QP and the second QP are part of a particular node.

4. The communications system according to claim 1,

wherein an incoming message is received by the first QP,

wherein at least one of the first QP and the resource manager determines a number of resources to request for the first QP, and

wherein at least one of the first QP and the resource manager determines whether the first QP is allowed to draw any resources from the SRQ.

5. The communications system according to claim 4, wherein at least one of the first QP and the resource manager checks whether the first limit value is zero or whether the first limit value is smaller than a number requested by at least one of the first QP and the resource manager.

6. The communications system according to claim 4, wherein the incoming message received by the first QP is dropped if the first QP is not allowed to draw any resources or enough resources from the SRQ.

7. The communications system according to claim 4, wherein the first connection is dropped if the first QP is not allowed to draw any resources or enough resources from the SRQ.

8. The communications system according to claim 4,

wherein the first QP draws resources from the SRQ, and

wherein at least one of the first QP and the resource manager decrements the first limit value based upon the resources drawn from the SRQ by the first QP.

9. The communications system according to claim 1,

wherein the general pool comprises a shared completion queue (SCQ),

wherein at least one of the first QP, the second QP and the resource manager monitors completions posted on the SCQ, and

wherein at least one of the first QP, the second QP and the resource manager adjusts the first limit value or the second limit value based upon the monitored completions posted on the SCQ.

10. The communications system according to claim 1,

wherein the first QP comprises a first completion queue (CQ),

wherein at least one of the first QP and the resource manager monitors completions posted on the first CQ, and

wherein at least one of the first QP and the resource manager adjusts the first limit value based upon the monitored completions posted on the first CQ.

11. A method for communications, comprising:

establishing a first connection associated with a first queue pair (QP);

establishing a second connection associated with a second QP;

sharing a shared receive queue (SRQ) among the first QP and the second QP;

assigning a local limit value to be associated with the first connection, the local limit value being associated with a maximum number of SRQ buffers that are accessible to the first QP; and

dropping the first connection or an incoming message on the first connection if the first QP attempts to access more SRQ buffers than its maximum number.

12. The method according to claim 11, wherein the local limit value associated with the first connection is managed by the first QP.

13. The method according to claim 11, further comprising:

receiving the incoming message via the first QP;

drawing one or more buffers from the SRQ; and

adjusting the local limit value based upon the one or more drawn buffers.

14. The method according to claim 13, wherein adjusting the local limit value comprises decrementing, by the first QP, the local limit value based upon the one ore more drawn buffers.

15. The method according to claim 11, further comprising:

monitoring one or more completions posted on a first completion queue (CQ) of the first QP; and

adjusting the local limit value based upon the one or more monitored completions.

16. The method according to claim 15, wherein adjusting the local limit value comprises incrementing, by the first QP, the local limit value based upon the one or more monitored completions.

17. The method according to claim 11, further comprising:

monitoring completions posted on a shared completion queue (SCQ), the SCQ being shared by the first QP and the second QP; and

adjusting the local limit value based upon the monitored completions.

18. A communications system, comprising:

a first queue pair (QP) associated with a first connection;

a second QP associated with a second connection;

a shared receive queue (SRQ) shared by the first QP and the second QP; and

a watermark indicating a low level of work queue elements (WQEs) on the SRQ.

19. The communications system according to claim 18, wherein, if a number of WQEs on the SRQ drops below the watermark, then an asynchronous event is generated.

20. The communications system according to claim 19, wherein the asynchronous event notifies at least one of an upper layer protocol (ULP) layer, a buffer manager, and any part of a consumer that the number of WQEs on the SRQ has dropped below the watermark.

21. The communications system according to claim 19, wherein at least one of the ULP layer, a buffer manager, and any part of a consumer analyzes completions posted on a shared completion queue (SCQ) shared by the first QP and the second QP.

22. The communications system according to claim 21, wherein the ULP layer determines which of the first QP or the second QP is transporting excessive traffic or offending traffic.

23. The communications system according to claim 18, wherein an out-of-order WQE limit is associated with the SRQ.

24. The communications system according to claim 18,

wherein a first out-of-order WQE limit is associated with the first QP, and

wherein a second out-of-order WQE limit is associated with the second QP.

25. A method for communications, comprising:

sharing a receive queue (RQ) between a first queue pair (QP) associated with a first connection and a second QP associated with a second connection;

setting a threshold that is associated with the shared RQ and is indicative of a particular number of work queue elements (WQEs) on the shared RQ; and

generating an asynchronous event if the threshold is reached.

26. The method according to claim 25, comprising:

analyzing completions on a completion queue (CQ) to determine if either the first QP or the second QP is draining the WQEs on the shared RQ.

27. The method according to claim 26, wherein the CQ comprises at least one of a dedicated CQ and a shared completion queue (SCQ).

28. The method according to claim 25, comprising:

setting an out-of-order WQE limit for the shared RQ.

29. The method according to claim 25, comprising:

setting a first out-of-order WQE limit for the first QP.

30. The method according to claim 25, comprising:

setting an out-of-order WQE limit which is part of another local limit.

31. The method according to claim 25, comprising:

dropping a WQE holding an out-of-order message from the SRQ upon reaching a particular watermark.